

Digital Storage Considerations and Devices for Archiving Motion Picture Assets

1. Introduction

Motion picture productions generate vast amounts of digital data. This data can be in the form of text, images, audio, video, and other forms (including metadata) generated during the pre-production stage – involving script generation, pre-visualization and virtual assets, post-visualization and production assets, to post-production – involving visual effects, editorial and sound design and final mastering assets, to the numerous deliverables for exhibition and distribution. With image acquisition modalities of 4K and 8K, stereo, high frame rate, as well as audio acquisition of large sound fields, on-set LiDAR data, and so on it is easy to see that a motion picture production's storage needs are very large.

Content owners and the motion picture industry have been archiving and preserving these assets for decades, utilizing processes such as transferring the images and audio tracks to film, or using magneto-optical tape for audio soundtracks (see <https://academydigitalpreservationforum.org/2024/07/26/understanding-the-forgotten-world-of-analog-film-sound-workflow-to-improve-sound-preservation-and-restoration/>). With the advent of digital tools in the production process, a common approach for archiving employed by most content owners is to store motion picture assets digitally on data storage devices. These storage devices are commonly large hard disk drives or tape-based storage, with LTO tape being most common, and predominantly retained on-premise. Data storage in the cloud, either public, i.e., storing data on shared infrastructure offered by cloud service providers or private, i.e., storing data on private infrastructure maintained either by the content owner or in a co-located data center, is an alternative that has become increasingly compelling in recent years. This paper examines the various considerations for storing motion picture assets in the cloud.

In the sections below, the paper examines the access, retention and security considerations for any digital archival storage, and looks at how tape and hard disk drive satisfy these considerations.

In a separate paper, *File Systems and the Cloud for Digital Archival of Motion Picture Assets*, data storage is explored at an enterprise level and data access is examined at that scale. That paper reviews file systems that access the storage blocks highlighted here and introduces distributed file systems and builds up to the capabilities offered by cloud providers today and the value-added scenarios offered by the public cloud providers that could benefit content owners.

2. Data Considerations

As we look at mechanisms for storing digital motion picture data, it is important to consider the different requirements for our data. We can then examine the different storage strategies and how they satisfy these requirements.

2.1. Data Retention

The data needs to be preserved for a long time or until it is explicitly deleted. The fidelity of the data needs to be maintained without any compromise, irrespective of whether it is short-lived or long-lived. It is important to note that the data needs to be preserved even in the face of catastrophes such as natural disasters.

2.2. Data Access

The data needs to be accessible and data access requirements vary across different types of assets. These requirements can be classified as low and high latency, which typically applies to the finished material and to the underlying production assets modified to make the finished material, respectively.

Also, data must be accessible over time, meaning it is stored on media supported and operable as technology changes. Every few years, technological advancements lead to new media (tape formats, file systems). For the stored data to be accessible, it then becomes important for the media to be backward compatible, or for media to be migrated to the new media or even a new format

2.3. Data Protection and Security

The data needs to be protected. The data needs to be protected from corruption and failure of the storage media, but also from any disaster that might strike the storage location – rendering access to the storage impossible, and thus loss of the archive.

The data needs to be secure, either from unauthorized access by bad actors, or improper or erroneous access by authorized users. And in fact, cloud service providers have implemented their own robust security protocols. But in some cases, it may be important for the data to be encrypted, so that if by some chance, access is obtained to the storage media, an additional layer of protection is available before the data is stolen. Of course, if data is additionally encrypted, there will also be the need for a long-term strategy for continued maintenance of the security keys to decrypt. For this reason, many archivists have expressed concerns about this additional encryption layer and instead focus on the inherent security models of their cloud vendors.

2.4. Data Utilization

The storage lifecycle of content for motion pictures typically starts with data that is minimal (at concept stage), grows to very large (at capture stage) and systematically reduces in a deliberate process (during editorial) to create a highly specific viewing experience (VFX and color correction) and then rapidly grows again to support all the distribution channels and multiple languages. From this it is evident that depending on the stage of the motion picture, while data is important to retain, not all the data is archival, and not all of the datasets are equally important in an archival context.

Data utilization should be viewed in the context of what stage a production is at, and the access, security, protection and storage needs for that stage. The storage architecture for that stage should then be constructed in a fashion that optimally supports those needs. While each production is a unique snowflake, for the purposes of discussion, we can view a project progressing through nine stages: pre-production, production, review, editorial, visual effects, color correction, localization, distribution and archival. These stages are neither linear nor clear-cut in their division, and depending on the production, shots will transition back and forth between stages, and depending on the kind of production tools used, such as virtual production, these boundaries will be nonexistent. However, to illustrate the notion of data utilization, we can begin with these nine stages.

In pre-production, the activity is limited to a small group, but they need to iterate rapidly, and the dataset is fairly small. (Note: This example does not consider the virtual production scenario). In production, we are dealing with large amounts of sound and video capture, and this activity is temporary and requires complex and rapid access in a secured environment. Then captured files have to be made available for review (as a smaller dataset, but to a larger group) as well as made available for collective access to editorial. The data pool is reduced in editorial, but as shots are provided to visual effects, the rendering needs increase the dataset and access requirements dramatically. As we approach finishing and mastering, the audio mixing and color correction bring the final data elements as high-resolution elements that require large amounts of data storage but the access is limited to a small group. This access pool has widened as many operators are required to generate all of the multiple-language versions and the dubbed versions, albeit with smaller proxy files. In distribution, the activity is focused on funneling data to downstream clients or distribution partners, requiring large access and large datasets and similar to production data, security, speed and structured, automated access are required characteristics. Lastly, the data can be archived, and this requires upload of the archival data set to long-term storage.

2.5. Other Considerations

In addition to the basic data considerations that are listed above, there are numerous enterprise-level requirements that need to be satisfied for any digital data archive. Some of these are:

1. Common ontologies and metadata in formats that are flexible, extensible and robust are important to support discoverability and automation. Similarly, file formats of files destined for archive should ideally not be in proprietary formats.

It is important that the production dataset's ontology is maintained, so that entity relationships within the dataset are available and accessible to anyone who accesses the digital archive. This ontology must be flexible, extensible (as per different production needs), robust and able to connect and map to other evolving and extending ontologies.

The metadata needs for a motion picture are extensive and vary drastically based on the production group that is working on it. For example, the on-set metadata requirements are needed by on-set personnel, and in some limited fashion by downstream departments like

editorial and visual effects, who in turn will add additional metadata. This metadata needs to be filtered appropriately for the group that needs it, yet at the same time needs to be flexible, extensible, robust, and easily transportable.

2. The repository needs to provide a secure mechanism so any authorized application can access assets via a common search and discovery interface.

Multiple levels of access control for users, administrators and operators in a secure fashion with continuous logging of access and activity is vital for any repository that stores the “jewels” of the motion picture. It is also necessary to provide this access via Application Programmable Interfaces (APIs) and to provide an easy to understand and navigable User Interface (UI/UX).

3. Management of assets is driven by policies that are enforceable regardless of any infrastructure on which they are hosted or the applications that manage them.
4. And many other considerations that are related to
 - a) Business Processes
 - b) Corporate Governance
 - c) Clearance, Rights and Legal issues relating to ownership
 - d) Etc.

3. Common Types of Storage for Digital Data

There are three distinct types of digital storage media that are commercially deployed and supported worldwide, namely, **magnetic Tape**, **magnetic Hard drive** and the **Solid-state drive**.

There are numerous other technologies and a review of those is outside the scope of this paper. The Academy Digital Preservation Forum invites proponents of other technologies to propose how a commercially viable and supportable system would satisfy the considerations outlined in Section 1.

3.1. Magnetic Tape

The first commercial digital-tape storage system, IBM’s Model 726, introduced in 1952, could store about 1.1 megabytes on one reel of tape. Today, a modern tape cartridge can hold 15 terabytes. And a single robotic tape library can contain up to 278 petabytes of data.

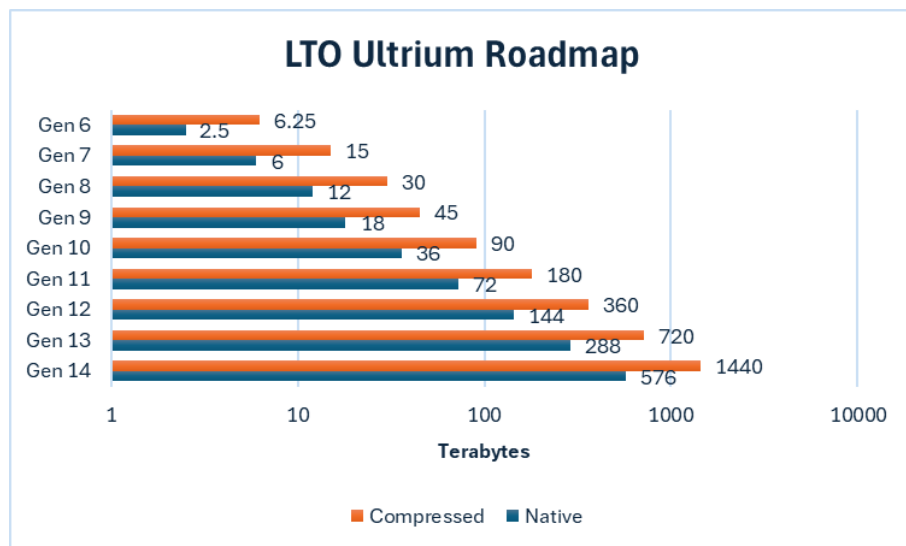
Note: *This paper specifically does not address the numerous tape formats used for broadcast audio and video in the Media and Entertainment industry. While there is precedent for archiving these tapes, the archival methodologies are predominantly like archiving physical assets; i.e., placed in a vault with temperature and humidity control, and outside the scope of this paper.*

Data is written to tape in the form of narrow tracks in a thin film of magnetic material in which the magnetism switches between two states of polarity. The information is encoded as a series of bits, represented by the presence or absence of a magnetic-polarity transition at specific points along a track. This information is written in blocks with inter-block gaps between them, and each block is written in a single operation with the tape running continuously during the write. The linear method of writing data (as opposed to the helical) arranges data in long parallel tracks that span the length of the tape. Multiple tape heads simultaneously write parallel tape tracks on a single medium.

The current state of the art (at the time of writing this paper, August 2024) for tape digital data storage is primarily Linear Tape Open (LTO) technology. There are a few instances of digital archives that store data on the IBM Jaguar Enterprise class of tapes and drives (current generation is known as TS11xx). Digital archiving (for M&E) has been carried out on tape since the early 1980s, with tapes such as the Digital Linear Tape (DLT), Ampex DST, Mammoth, AIT, StorageTek (Redwood and T9xxx) and many other proprietary and now obsolete formats. These formats and drives are not discussed in this paper; suffice to say that any data on these tapes should be migrated onto the new generation of tapes as soon as possible.

With the introduction of LTO, the industry has leveraged, since the early 2000s, an open format to enable interoperability between drives and media from multiple licensed vendors. This has permitted content owners to have the freedom to choose from a combination of drives, libraries and cartridges without the risk of being locked into proprietary technology. With a published roadmap for over 14 generations, new generations have been introduced regularly by the LTO consortium with higher capacity and transfer rates.

With backward compatibility (up to Gen 7) permitting write back one generation and read back two generations, the LTO consortium has enabled time for digital archivists to migrate from an older tape generation to a newer one. As an example, an LTO-7 drive will read/write LTO-6 and read LTO-5. Note that for the later generations (after 7, i.e., 8th and 9th gen), LTO Ultrium tape drives are ONLY able to write back and read back one generation.



As an example, an LTO-7 drive will read/write LTO-6 and read LTO-5. Note that for the later generations (after 7, i.e., 8th and 9th gen), LTO Ultrium tape drives are ONLY able to write back and read back one generation.

As seen in Figure 1, the LTO consortium has published a clear roadmap for LTO generations, and each new generation increases storage capability and at the same time improves data access rates. So, for LTO 9, the sustained data r/w rate is 50 MBs, and for LTO 8 the sustained data r/w rate is 45 MBs.

LTO tape also addressed, with the introduction of LTFS (Linear Tape File System), a major drawback of magnetic tape drives, namely, the ability to store metadata. Magnetic tape does not hold file metadata independent of the file content data on tape, as data is written to tape in blocks with inter-block gaps between them, and each block is written in a single operation with the tape running continuously during the write. A simple index at the head of the tape is used to maintain a notion of the file name, timestamps, and directory hierarchy, and the rest of the metadata storage is delegated to external databases that a user must custom develop and separately maintain. For this generation of archives the UNIX 'tar' interoperable standard was used for storing data on tape; however, the 'tar' utility did not permit modification of file metadata independent of modifying file content data, and did not provide an overall view of the files stored on the tape or their characteristics.

LTFS was developed and specified as an open ISO standard and introduced in LTO drives in 2010. LTFS enables a partitioning feature that helps to enhance file control and space management. Partitioning allows for a portion of the tape to be reserved for indexing, which tells the drive precisely where in the tape a file is stored. This second partition holds the actual file. With LTFS, the indexing information is first read by the drive and presented in a simple, easy-to-use format that allows for usage like a hard drive. The index portion of the tape permits storage of metadata about the file and thus can present a detailed file hierarchy tree that all of the attribute's users are familiar with in a desktop file browser.

Most manufacturers estimate that LTO tapes will provide dependable data storage for 15-30 years. However, LTO archives require precise storage conditions. Slight deviations in temperature or humidity can affect storage life. Ideally, LTOs should be stored at 61 to 77°F (16 to 25°C).

Error handling for tapes is carried out at two levels,

- a) the way in which the data is written to the tape (Error Correction Code) and,
- b) the way in which the tape behaves within the drive.

At the first level, LTO uses Error Correction Code (ECC), at two levels (C1 and C2) to protect data written to the tape. C1 and C2 are highly efficient Reed-Solomon ECC. C1 is designed to detect and correct media errors, and C2 operates as an Erasure code to correct any residual errors from the C1 process. These codes are added to the data before it is written to the tape and then verified after it is written, to ensure an extremely low-bit error rate (BER). Typical operating numbers for BER are 1 error event in 10¹⁹ bits, and to put that number in context, it means that an LTO tape would encounter one error bit in an Exabyte (1000 Petabytes) that could not be fixed by the error correction technology employed by LTO.

At the second level, the tape drive also verifies the tape as it reads and writes to the tape. It records an error count, for both reads and writes, onto a memory chip inside the tape cartridge. When the tape is loaded, enterprise-level tape drives read this error count, and if it exceeds a particular threshold, permit corrective action. The error count cannot be reset so the history of previous errors is never lost.

When writing data to a tape, if an error is encountered, the drive simply retries and rewrites the data further down the tape. This is logged as a recovered write error and recorded onto the memory chip. While low recovered write errors are no problem, large counts can indicate poor tape quality or a drive issue. Recovered read errors indicate that the drive needed to retry the read; again low counts are normal. But high rates in the hundreds or thousands indicate a problem with the tape or drive, and the ideal recourse is to migrate the data onto another tape, on another drive.

While LTO error handling is robust, it alone does not tell the complete picture. The actual error rate is a combination of real-world performance, tape wear, drive wear, etc. and while there are anecdotal figures on tape replacement due to errors, no published studies are available to reference on the real-world error rate of LTO tapes and drives. Anecdotally, the author's experience has been that across 8000 tapes, two tapes had to be replaced over a four-year period – resulting in an annualized failure rate (AFR) of 0.00625%.

3.2. Hard Disk Drive (HDD)

Recognizing a major drawback of tape storage systems, i.e., one of serially accessing the data, IBM in 1953 invented the hard disk drive – with the promise of high capacity, rapid random access at lower cost. From an original disk capacity of 3.75 MB with a footprint of 24 inches, the technology has evolved to currently providing 30 TB of storage in 2.5 inches.

HDDs record data by instantiating sequential changes of magnetism in a thin film of ferromagnetic material on both sides of a disk, with the direction of magnetization representing binary data bits. The data is read from the disk by detecting the transitions in magnetization. For writing efficiency, data is encoded using a run-length limited encoding scheme.

Structurally, a typical HDD consists of a spindle that holds flat circular disks, called platters, which hold the recorded data. The platters are coated with a shallow layer of magnetic material typically 10-20 nm in depth, about




5000 times thinner than a piece of paper. The platters are rotated at high speed (from 4200 rpm in energy-efficient portable devices, to 15,000 rpm for high-performance servers) past read-and-write heads that are positioned very close to the magnetic surface. With a distance of tens of nanometers, these read-and-write heads detect and modify the magnetization of the material passing immediately under it. With one head for each magnetic platter surface on the spindle, mounted on a common arm, access to the entire surface of the platter is carried out by an actuator arm that moves the heads in an arc across the platters as they spin.

As illustrated in Figure 2, the HDD has 2 moving parts: A spindle motor that spins the disks and an actuator motor that positions the read/write head assembly across the spinning disks. These moving parts are the weakest link in the HDD solution, resulting in a service lifetime of between 3-5 years. A common way of prolonging the lifespan of an HDD is to spin it down, meaning power down the drive to stop the spindle and actuator motors from spinning. However, when not in use, HDDs can suffer from issues like stiction, where the read/write head gets stuck to the platter, or lubricant in the bearings dries out.

Hard drive manufacturing (and manufacturers) peaked in the early 1980s and since then the industry has consolidated around three major manufactures, namely Western Digital, Seagate, and Toshiba. The number of HDD units peaked in 2010 at about 650 million units and has steadily declined since then, with approximately 166 million units shipped in 2022.

HDDs use error correction codes (ECCs) for error handling. While read/write heads are extremely accurate, they occasionally “flip” a charge. In other words, the heads fail to write a charge or write a charge incorrectly. Without error correction, data corruption would be extremely common. To prevent this each sector of a hard drive includes information about the data, which allows incorrect data to be corrected by the hard drive controller before it is transmitted to the computer. Most modern hard drives use the Reed-Solomon error correction algorithm, and the level of redundancy directly corresponds to the number of ECC bits included in each sector. When the hard drive reads data, the ECC informs the controller whether the data contains errors. If the number of errors is within the tolerable range of the ECC, they can be fixed. Enterprise HDDs are rated as one in 10^{14} bits read, and to put that number in context, it means that an enterprise level HDD would encounter one error bit in 12.5 Terabytes that could not be fixed by the error correction technology employed by HDDs.



Despite the error correction technology employed by HDDs, the actual failure rates of hard drives are considerably higher. The company Blackblaze provides a report on the annualized failure rates of a large number of hard drives. This report is accessible here (<https://www.backblaze.com/blog/backblaze-drive-stats-for-2023/>).

The findings for 2023, by Backblaze, show that 270,000 hard drives, across 35 different hard drive models, running for a combined total of nearly 90 million days in 2023, reported an annualized failure rate (AFR) of 1.70%. To put this in perspective, this figure of 1.7% represents about 4,590 disks that failed across the 270,000 drives. Additionally, their calculations show that “on average” the expected lifetime of an HDD is about two years and seven months. In the company’s words, “Doing a little math, over the last year on average, we replaced a failed drive every two hours and five minutes. If we limit hours worked to 40 per week, then we replaced a failed drive every 30 minutes.”


3.3. Solid-State Disks (SSD)

Recognizing a major drawback of Hard Disk Drives, i.e., the moving parts of the spindle and the read/write head actuator, manufacturers introduced Solid State disks (SSDs) in the late 1970s. In 1978, the STC 4305 from StorageTek was offered as a direct plug-compatible replacement device for the IBM 2305 hard drive. This product provided 45 MB of storage for a cost of around \$400,000 for a 45 MB capacity but at a 7x higher data read/write rate of 1.5 MBs. Today (August 2024), SSDs can provide storage capacities that exceed 100 TB, at read/write rates of 15 GBs and a cost point of \$0.05/GB.

SSDs use non-volatile memory, typically NAND flash^{iv}, on integrated circuits to store data. Flash memory stores data in individual memory cells, which are made of floating-gate transistors. These integrated circuits retain their information when no power is applied to them.

Unlike HDDs, SSDs have no moving parts, allowing them to deliver faster data access speeds, reduced latency, increased resistance to physical shock, lower power consumption, and silent operation. SSDs are interfaced to the host computer in the same way as HDDs and are offered a direct replacement with the same size footprint. In addition, SSDs are also offered as hybrid storage solutions, such as solid-state hybrid drives (SSHDS), combining SSD and HDD technologies to offer improved performance at a lower cost than pure SSDs. SSD performance and endurance vary depending on the number of bits stored per cell, ranging from high-performing single-level cells (SLC) to more affordable but slower quad-level cells (QLC). In addition to flash-based SSDs, other technologies such as 3D XPoint (commercially offered as Intel’s OPTANE) provide faster speeds and higher endurance through different data storage mechanisms.

SSDs are generally more expensive on a per-gigabyte basis, yet SSDs are increasingly replacing HDDs, especially in performance-critical applications and as primary storage in many consumer devices. SSDs, however, do have limitations, such as data loss over time, a



finite number of write cycles, and repeated read/write causing error loss.

SSDs based on NAND flash slowly leak charge when not powered, while heavily used consumer drives may start losing data typically after one to two years in storage. Additionally, due to the way data is written and read to the drive, SSD cells can wear out over time. The process of pushing electrons through a gate to set its state wears on the cell and over time reduces its performance until the SSD wears out. While this effect can take a long time, and there are mechanisms to minimize this effect (e.g., TRIM), it is important to keep in mind while selecting this device as a storage medium. Also, SSDs have a limited lifetime number of writes, and slow down as they reach their full storage capacity.

Like HDDs, SSDs also use ECCs for error correction, although different error correction algorithms (such as the BCH or Hamming algorithms) are utilized. When data is written to an SSD, ECC codes are generated based on the content of each data block. These ECC codes are then stored alongside the data. During data retrieval, the ECC codes are recalculated, and any discrepancies indicate the presence of errors. The ECC algorithms utilize parity bits or more complex mathematical calculations to identify and correct errors automatically. Enterprise SSDs are rated as one in 10^{-16} bits read, and to put that number in context, it means that an enterprise level SSD would encounter one error bit in a Petabyte (1000 Terabytes) that could not be fixed by the error correction technology employed by SSDs.

In spite of the error correction technology employed by SSDs, the actual failure rate in the field is considerably higher. The company Blackblaze provides a report on the annualized failure rates of many solid-state drives. This report is accessible here (<https://www.backblaze.com/blog/ssd-edition-2023-mid-year-drive-stats-review/>).

The findings by Backblaze, mid-year for 2023, show that 3,144 SSDs, across 14 different SSD models, running for a combined total of nearly 277 thousand days, reported an annualized failure rate of 1.05%. To put this in perspective, this figure of 1.05% represents about 33 disks that failed across the 3,144 drives. In the company's words, "we just don't have enough data to get decent results. For any given drive model, we like to see at least 100 drives and 10,000 drive days in a given quarter as a minimum before we begin to consider the calculated AFR to be "reasonable."

3.4. Tape vs HDD vs SSD

How these devices satisfy the needs from Section 2

	Tape	HDD	SSD
Data Retention	<ul style="list-style-type: none"> 15-30 years. Ideally, stored at 61 to 77°F (16 to 25°C). 	<ul style="list-style-type: none"> 2-3 yrs. 	<ul style="list-style-type: none"> 1-2 yrs.
Data Access	<ul style="list-style-type: none"> 45-50 MBs Migrate data every 3-5 years With LTFS, an open protocol 	<ul style="list-style-type: none"> 100-200 MBs Migrate data every 2-3 years Most HDDs have a standard block size and data can be accessed independent of OS and controller. 	<ul style="list-style-type: none"> 300-500 MBs Migrate data every 1 year SSDs come in a variety of block and memory architectures, and while some are standard, there are proprietary ones too.
Data Protection and Security	<ul style="list-style-type: none"> With built-in, on-the-fly encryption and additional security provided by the nature of the medium itself, tape offers inherent security. Additionally, if a cartridge isn't mounted in a tape-drive, the data cannot be accessed or modified. This "air gap" is particularly attractive from a security standpoint, and can prevent inadvertent access or worse, inadvertent deletion. 	<ul style="list-style-type: none"> Enterprise-level HDDs offer protection for user data at rest, through encryption and access control, to ensure that the data can only be accessed by authorized users in the case of loss, theft, or remote attack. 	
Notes	<ul style="list-style-type: none"> Very energy efficient: A tape cartridge simply sits quietly in a slot in a robotic library and doesn't consume any power at all. 	<ul style="list-style-type: none"> HDDs can be spun down for energy efficiency. 	

As seen from the table, all of the devices can satisfy the requirements called out in Section 2, with some limitations. The data requirements called out in Section 2.4 are enterprise-level requirements and are typically supported at a file system level or with external systems that augment the storage units.

4. Additional Strategies for Digital Archival

As described in Section 3 above, the storage media of tape, HDD and SSD are not infallible, and in fact are destined to fail by the very nature of their design. While the bit error rate (BER) advertised by these technologies and their error correction methodologies are promising, the actual failure rate (AFR) from in-field testing tells an entirely different story. So, what is a digital archivist to do? They must employ additional strategies of increasing the data integrity and adding data redundancy.

4.1. Increase Data Integrity

4.1.1. *Scrubbing*

To mitigate adjacent-track interference that occurs when bits encoded on one track corrupt bits already written on adjacent tracks the data is periodically read and re-written. This technique is called scrubbing. It helps in reducing the probability of an error when reading data from the disk. Many commercially available enterprise file systems perform scrubbing every two weeks on average.

4.1.2. *Checksumming*

To mitigate bit flips or data corruption during transfer or at rest, a checksum is computed on the data, and this checksum is recomputed at every stage of transit or when the data is at rest as a verification process. Additionally, it is recommended that periodic integrity checks are performed at the storage media level, by reading the stored data, recomputing the checksum and validating it against the previous checksum that was computed when the data was written to the media.

4.2. Add Data Redundancy

4.2.1. *Add Parity*

Parity information refers to the additional data stored in a physical array to provide data protection in case of either a tape, HDD or SSD failure. It allows for the implementation of a lossless bitwise exclusive-OR algorithm to replace the missing data and maintain data integrity. There are numerous parity schemes, utilized in a redundant array of independent disks (RAID), and is not just limited to disks, but can also be deployed for tape. Depending on the required level of redundancy and performance, different schemes, or data distribution layouts, named by the word “RAID” followed by a number, for example RAID 0 or RAID 1 can be deployed. Each RAID level (1-6) provides a different balance among the key goals: reliability, availability, performance, and capacity.

4.2.2. Replication

Replication is a simple approach to adding redundancy to data, which involves storing multiple copies of data on separate devices. For example, the three-replication scheme stores three copies of data on three separate devices.

4.2.3. Geographical redundancy

No amount of local replication will protect against circumstances that affect the storage location or data center. For example, natural calamities such as earthquakes and severe storms could end up affecting entire data centers. Protection against such catastrophic events is to add redundancy across geographical locations. In such scenarios, the replicas of the data are stored in different geographical locations.

5. Conclusion

In the sections above, we examined the access, retention and security considerations for any digital archival storage, and looked at how tape, hard disk drive and solid-state disk drive satisfy these considerations at a unit level. We also reviewed the data utilization for motion pictures, and as all of the devices satisfy the requirements they are deployed for storing and archiving motion picture assets.

To be able to address the storage needs of a motion picture, it is necessary to pool multiple units of Tape, HDD and SSD together. There are numerous strategies for combining multiple devices, and in another paper, *File Systems and the Cloud for Digital Archival of Motion Picture Assets*, we will explore the advantages of a file system, distributed file systems and servers, and distributed computing. At the device level, this paper illustrates how digital archival can be carried out on Tape, HDD and SSD. Depending on the budget, retention needs as well as access and retrieval requirements, most digital archivists will use a combination of these devices in their archival strategy.

It should be pointed out that all of the mechanisms described here have finite life and device failure is expected and should be planned for in any archival strategy. This can be in the form of redundancy, i.e., multiple copies, geographical separation, and regular maintenance and verification. Also, as with all technical products, an archivist must mitigate obsolescence, and the only way is to migrate the content from older (and prone to failure) devices onto newer devices. This migration step is also the perfect time to validate the data and ensure that there has been no data corruption.

While digital storage devices are not write-once and store devices (à la film), it should be pointed out that they are commercially supported and service industries outside of Media and Entertainment. This implies that support for these devices will be widely available and will continue to improve. Further, digital tools are ubiquitous and are commonly used, thus it is possible to access any data repository (given the right user authorization) without the need for custom hardware or specialized software – in fact, most times a simple web browser will suffice.

Such democratization of tools and devices makes them easy and cheap to use, as well as widely adopted, leading to great support in the future.

6. References

- i. <https://academydigitalpreservationforum.org/2024/07/26/understanding-the-forgotten-world-of-analog-film-sound-workflow-to-improve-sound-preservation-and-restoration/>
- ii. <https://spectrum.ieee.org/why-the-future-of-data-storage-is-still-magnetic-tape>
- iii. <https://spectrum.ieee.org/tape-is-back-and-better-than-ever>
- iv. <https://www.techtarget.com/searchstorage/definition/NAND-flash-memory#:~:text=NAND%20flash%20memory%20is%20a,devices%2C%20such%20as%20hard%20disks.>
- v. <https://brilliant.org/wiki/error-correcting-codes/>
- vi. <https://www.statista.com/statistics/398951/global-shipment-figures-for-hard-disk-drives/>